arXiv:1611.06492v1 [cs.CV] 20 Nov 2016

Recurrent Memory Addressing for describing videos

Kumar Krishna Agrawal Arnav Kumar Jain Abhinav Agarwalla Pabitra Mitra
Indian Institute of Technology Kharagpur
{kumarkrishna, arnavkj95, abhinavagarawalla, pabitra}@iitkgp.ac.in

Abstract

Deep Neural Network architectures with external memory components allow the model to perform inference and capture long term dependencies, by storing information explicitly. In this paper, we generalize Key-Value Memory Networks to a multimodal setting, introducing a novel key-addressing mechanism to deal with sequence-to-sequence models.

The advantages of the framework are demonstrated on the task of video captioning, i.e generating natural language descriptions for videos. Conditioning on the previous time-step attention distributions for the key-value memory slots, we introduce a temporal structure in the memory addressing schema. The proposed model naturally decomposes the problem of video captioning into vision and language segments, dealing with them as key-value pairs. More specifically, we learn a semantic embedding (v) corresponding to each frame (k) in the video, thereby creating (k, v) memory slots. This allows us to exploit the temporal dependencies at multiple hierarchies (in the recurrent key-addressing; and in the language decoder). Exploiting this flexibility of the framework, we additionally capture spatial dependencies while mapping from the visual to semantic embedding. Extensive experiments on the Youtube2Text dataset demonstrate usefulness of recurrent key-addressing, while achieving competitive scores on BLEU@4, METEOR metrics against state-of-the-art models.

1. Introduction

Generating natural language descriptions for images and videos is a long-standing problem, in the intersection of computer vision and natural language processing. Solving the problem requires developing powerful models capable of extracting visual information about various objects in an image, while deriving semantic relationships between them in natural language. For video captioning, the models are additionally required to find compact representations of the video which capture the temporal dynamics across image frames.

Figure 1. Our model generates employs a temporal attention mechanism on the visual features to identify key frames in the video. These are mapped to semantic features in the language domain for better context to the language model, which then generates the output sequence. Previously attended frames, and generated words identify the key frames, for generating the next word.

The recent advances in training deep neural architectures have significantly improved in the state-of-the-art across computer vision and natural language understanding. With impressive results in object detection and scene understanding, Convolution Neural Networks (CNNs) [23] have become the staple for extracting feature representations from images. Recurrent Neural Networks (RNNs) with Long Short Term Memory (LSTM) [15] units or Gated Recurrent Units (GRUs)[10], have similarly emerged as generative models of choice for dealing with sequences in domains ranging from language modeling, machine translation to speech recognition. Advancements in these fundamental problems make tackling challenging problems, like captioning [16, 45], dialogue [32] and visual question answering [1] more viable.

Despite the fundamental complexities of these problems, there has been increasing interest in solving them. A common underlying approach in these proposed models is the notion of ”attention mechanisms”, which refers to selectively focusing on segments of sequences [3, 47] or images [35] to generate corresponding outputs. Such attention based approaches are specially attractive for captioning problems, since they allow the network to focus on patches of the image conditioned on the previously generated tokens [45, 16], often referred to as spatial attention.

Models with spatial attention, however cannot be readily used for video description. For instance, in the Youtube2Text dataset, a video clip stretches around 10 sec-

onds, or around 150 frames. Applying attention on patches in these individual frames provides the network with local spatial context. This however, does not take ordering of the frame sequence or events ranging across frames, into consideration. To incorporate this *temporal attention* into the model, [47, 49, 27] extend this *soft alignment* to video captioning. Most of these approaches, treat the problem of video captioning in the sequence-to-sequence paradigm [37] with attentive encoders and decoders. This requires them to find a compact representation of the video, which is passed as context to the RNN decoder.

However, we identify two primary issues with these approaches. For one, applying attention sequentially provides the model with local context at the generative decoder [46]. As a result the decoder would be unable to deal with long-term dependencies while captioning videos of longer duration. Secondly, these models jointly learn the multimodal embedding in a visual-semantic space [28, 49] at the RNN decoder. With the annotated sentences being the only supervisory signal, learning a mapping from a sequence of images to a sequence of words, is difficult. This is specially true for dealing with video sequences, as the underlying probability distribution is distinctively multimodal. While [28] tries to address this issue with an auxiliary loss, the model suffers from the first drawback.

To address the aforementioned issues, we introduce a model which generalizes Key-Value Memory Networks [25] to a multimodal setting for video captioning. The proposed model, naturally tackles the problem of maintaining long-term temporal dependencies in videos, by explicitly providing all image frames for selection at each time step. We also propose a novel key-addressing scheme (see Section 4), which keeps track of previous attention distributions, allowing us to exploit temporal structure at multiple hierarchies and provide a global context. At the same time, the framework provides an effective way to deal with the complex transformation from the visual to language domain. Using a pre-trained model to explicitly transform individual frames (*keys*) to semantic embedding (*values*), we construct memory slots with each slot being a tuple (*key, value*). This allows to provide a weighted pooling of textual features as context to the decoder RNN, which is closer to the language model.

In summary, our key contributions are following :

- We generalize Key-Value Memory Networks (KV-MemNN) in a multimodal setting to generate natural descriptions for videos, and more generally deal with sequence-to-sequence models.

- We propose a novel key-addressing schema, conditioning on the previous time-step attention distributions for the key-value pairs. This allows us to exploit spatial and temporal structure at multiple hierarchies.

- The proposed model is evaluated on the YouTube dataset [7], where we outperform strong baselines while reporting competitive results against state-of-art models.

The remaining part of paper is organized as following. Section 2 provides a brief literature review of related work. In Section 3 we present details of our proposed model with recurrent key-addressing. Section 4 demonstrates the advantages of the model with experimental details and empirical analysis, followed by discussions and conclusions in Section 5.

2. Related Work

Following the success of end-to-end neural architectures and attention mechanisms, there is a growing body of literature for captioning tasks, in images and more recently videos. To deal with the multimodal nature of the problem, classical approaches relied on manually engineered templates [19, 11]. And while some recent approaches in this direction show promise [13], the models lack generalization to deal with complex scenes, videos.

As an alternative approach, [14, 18] suggest learning a joint visual-semantic embedding, effectively a mapping from the visual to language space. The motivation of our work is strongly aligned with [31], who generate semantic representations for images using CRF models, as context for the language decoder. However, our approach significantly differs in the essence that we capture spatio-temporal dynamics in videos while generating the text description.

Building on this, and Encoder-Decoder [3, 9] models for machine translation , [42, 41] develop models which compute average fixed-length representations (for images, videos respectively) from image features. These *context* vectors are provided at each time step to the decoder language model, for generating descriptions. The visual representations for the images are usually transferred from pre-trained convolution networks [33, 38].

An immediate drawback of the above approach is collapsing features across image frames to fixed vector by mean-pooling. For one, this loses the temporal structure across frames by treating them as "bag-of-images" model. Addressing this, [37] propose Sequence-to-Sequence models for accounting for the temporal structure, and [40] extend it to a video-captioning setting. However, passing a fixed vector as context at each time step, creates a bottleneck for the flow of gradients using Backpropagate Through Time (BPTT) [43] at the encoder.

Meanwhile, the notion of *visual attention*, has a rich literature in Psychology and Neuroscience, and has recently found application in Computer Vision [26] and Machine Translation [3]. Allowing the network to selectively focus on the patches of images or segments of the input se-

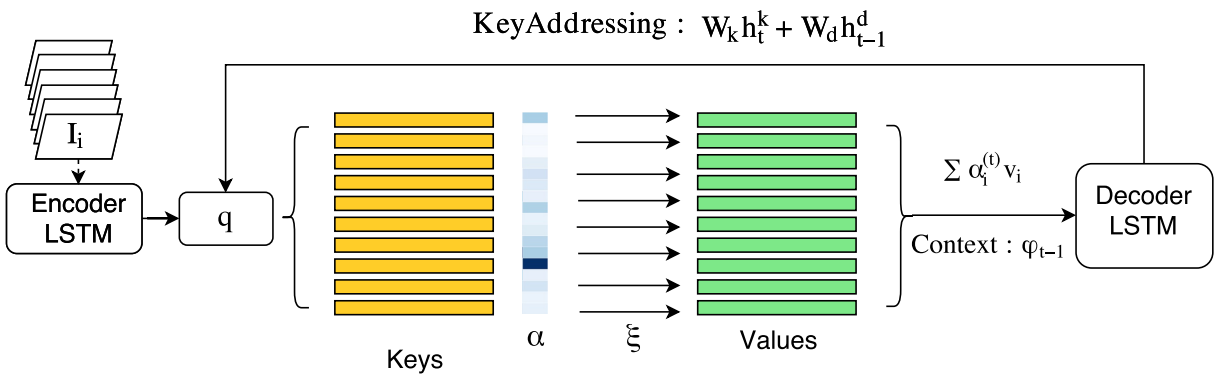

Figure 2. The video is considered as a sequence of image frames $\{I_1, ..., I_n\}$. For each key-value pair (k_i, v_i), the key k_i is features extracted from pretrained VGG-16 network, while the value v_i is text embedding extracted from pretrained image captioning network. The α_i correspond to the attention weights associated to the memory slots (k_i, v_i); $(h_{t-1}, h_t, ...)$ being the hidden states of the decoder RNN.

quences, representative works [45, 47, 16, 4, 49, 27, 28] have significantly pushed the state-of-the-art in individual domain. In essence, with the decoder output selectively conditioned on the encoder states the issues of fixed length representation and gradient bottleneck are largely solved:

$$p(y_i|y_{i-1}, ..., y_1, x) = g(y_{i-1}, s_i, c_i) \quad (1)$$

where,y_i is the readout, c_i is the context from the encoder and $s_i = f(s_{i-1}, y_{i-1}, c_i)$, is the hidden state of decoder RNN (See [3] for details).

However, as discussed in Section 1, sequential attention provides the decoder with local context [46]. Additionally, providing a semantic input which is closer to language space, as context to the decoder significantly improves on the capability of the model [48]. Our work closely brings these advances together in a Memory Networks framework [44, 36, 20]. While we introduce Key-Value Memory Networks [25] in a multimodal setting, there are several other key differences from previous works. For one, to our knowledge this is the first work which introduces video captioning in light of Memory Networks. This automatically deals with problems of maintaining long-term dependencies in the input stream by explicitly storing image representations. Meanwhile we also tackle the "vanishing gradient problem" typical with training RNN Encoder-Decoder modules for long input sequences.

Key-Value MemNNs [25] are originally proposed for QA task in the language domain, providing the last time-step hidden state, as input to the classifier. In this work, we address a more complex problem of video captioning by proposing a novel key-addressing scheme (details in Section 4) and *(key, value)* setup for exploiting the hierarchical spatio-temporal structures. The model tracks the attention distribution at previous time steps, thereby providing a strong context on where to attend on the complete video sequence. This implicitly provides a global temporal structure at each readout. While similar in motivation to [48, 30], the model architecture and domain of application, especially on capturing global temporal dynamics in videos as opposed to images or entailment, is significantly different.

3. Recurrent Memory Addressing for videos

Our model is a general extension of the encoder-decoder framework [9, 3, 45, 42, 17], in a Memory Networks [44, 25] setting. The encoder network learns a mapping from the input sequence to a fixed-length vector representation, which is used by the decoder to generate output sequences. Unlike standard Encoder-Decoder architectures, our model (see Fig. 2) comprises an encoder module, key-value memories and a decoder module.

3.1. Encoder

The encoder network E maps a given sequence $X = \{I_1, ..., I_T\}$ to the corresponding sequence of context representation vectors $\{k_1, ..., k_T\}$. As we are dealing with videos (sequence of images), we define two different encoders to achieve the mapping.

CNN Encoder: Given an input image $I_t \in \mathbb{R}^{NxM}$, the CNN encoders learn a mapping $f : \mathbb{R}^{NxM} \rightarrow \mathbb{R}^D$, if we consider the output of fully-connected layers in standard ConvNet architectures [33]. In case of convolution outputs, the encoders learn $f : \mathbb{R}^{NxM} \rightarrow \mathbb{R}^{LxD}$, where L is the number of context vectors of dimensionality D.

RNN Encoder: Consider an input sequence of length T. The RNN encoder processes the input X sequentially, generating hidden states h_t at each time step, where

$$h_t = g(f(I_t), g_{t-1}) \quad (2)$$

While maintaining temporal dependencies, this allows us to map variable length sequences to fixed length context vector. In this work we use modified version of an LSTM [15] unit, as proposed in [50] to implement g.

Learning a good representation for the *keys* is imperative to good performance of the model. We experiment with different settings of the encoder, using features directly from the CNN Encoders, or stacking an RNN Encoder on these extracted feature vectors.

3.2. Key-Value Memories

The model is built around a Key-Value Memory Network [25] with memory slots as vector pairs $(k_1, v_1), ..., (k_n, v_n)$. The keys and values serve the purpose of transforming visual space context into language space, and effectively capture the relationships between the video features and textual descriptions. The memory definition, addressing and reading schema is outlined below :

Keys (K): In our model, feature vectors for each frame I_t of the video are extracted using CNN Encoders. To incorporate sequential structure (video being a sequence of images), these appearance feature vectors are passed through a RNN Encoder, and hidden state at each timestep is extracted as key k_t.

$$k_t = g(f(I_t), k_{t-1}) \tag{3}$$

This allows the model to capture the temporal variation between frames, and take the ordering of actions and events into consideration. Implicitly this also helps, preserving high level information about motion in the video [4].

Values (V): Jointly learning visual-semantic embedding with supervisory signal only from annotated descriptions in Encoder-Decoder models is difficult [48]. To mitigate this, and to improve the context provided to the language model at decoder, we explicitly precompute semantic embedding corresponding to individual frames in the video.

For each image frame I_t we generate a semantic embedding v_t, which represents the *textual meaning* of that frame. This could be obtained from any pretrained model which jointly models visual and semantic embedding for images [45, 16, 42]. In this work we use a pretrained DenseCap model [16], to obtain semantic embedding corresponding to every image. Therefore,

$$v_i = \psi(I_t) \tag{4}$$

where, ψ is the intermediate visual-semantic embedding from a pretrained image captioning model. This gives for each frame in the video a key-value memory slot (k_t, v_t).

Key Hashing: This refers to retrieving a subset of the relevant memory slots for a given query vector q. For each video clip (typically with 120-240 frames), key-value pairs of T equally spaced image frames are extracted using pretrained models and are stored in the memory.

Key Addressing: This corresponds to the soft-attention mechanism deployed to assign a relevance probability α_i to each of the retrieved memory slots. These relevant probabilities are used for value reading (details in Section 4).

Value Reading: The value reading from the memory slots, returns the weighted sum of the key-value feature vectors: $\phi_t(K)$ and $\phi_t(V)$ at each time step. $\phi_t(K)$ is used for key addressing at the next time step and $\phi_t(V)$ is taken as input to the decoder RNN for generating the next word.

$$\phi_t(V) = \sum_{i=1}^{N} \alpha_i^{(t)} v_i \tag{5}$$

3.3. Decoder

The joint probability over generating the output sequence can be decomposed as:

$$p(y|x_1, ..., x_T) = \prod_{t=1}^{N} p(y_t|h_T, y_1, .., y_{t-1}) \tag{6}$$

where h_T is the context vector representing $X = \{x_1, ..., x_T\}$ and $y = \{y_1, ..., y_N\}$ is the output sequence. To train the network, the overall objective is therefore to maximize the log-likelihood of the output word sequence.

To this effect, Recurrent Neural Networks have been widely used for natural language generation tasks like Machine Translation, Image Captioning and Video Description generation. Vanilla RNNs are difficult to train for long range dependencies as they suffer from the *vanishing gradient problem* [5]. Thus, we use Long Short Term Memory(LSTM) [15] as they are known to memorize context for longer period of time using controllable memory units.

The LSTM model has a memory cell c_t in addition to the hidden state h_t of RNNs, which effectively summarizes the information observed up to that time step. There are primarily three gates which control the flow of information i.e (input, output, forget). The input gate i_t controlling the current input x_t, forget gate f_t adaptively allowing to forget old memory and output gate o_t deciding the extent of transfer of cell memory to hidden state. The recurrences at the decoder are defined as:

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{h}_{t-1} + \mathbf{U}_i \mathbf{x}_t + \mathbf{A}_i \phi_t(V) + \mathbf{b}_i) \tag{7}$$
$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{h}_{t-1} + \mathbf{U}_f \mathbf{x}_t + \mathbf{A}_f \phi_t(V) + \mathbf{b}_f) \tag{8}$$
$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{h}_{t-1} + \mathbf{U}_o \mathbf{x}_t + \mathbf{A}_o \phi_t(V) + \mathbf{b}_o) \tag{9}$$
$$\tilde{\mathbf{c}}_\mathbf{t} = \tanh(\mathbf{W}_c \mathbf{h}_{t-1} + \mathbf{U}_c \mathbf{x}_t + \mathbf{A}_c \phi_t(V) + \mathbf{b}_c) \tag{10}$$
$$\mathbf{c}_t = \mathbf{i}_t \odot \tilde{\mathbf{c}}_\mathbf{t} + \mathbf{f}_t \odot \mathbf{c}_{t-1} \tag{11}$$
$$\mathbf{h}_t = \mathbf{o}_t \odot \mathbf{c}_t \tag{12}$$

where $\odot$ is an element wise multiplication, σ is the sigmoidal non-linearity. $\mathbf{W}^i, \mathbf{U}^x, \mathbf{A}^x$ and $\mathbf{b}^x$, are the weight matrices for the previous hidden state, input, value context and bias respectively. Here $\phi_t(V)$ represents the context vector to the decoder at each time step.

Following standard sequence-to-sequence models with generative decoders, we apply a single layer network con-

ditioned on the hidden state h_t.

$$\mathbf{s}_t = \tanh(\mathbf{W}_s[h_t, x_t, \phi_t(V)] + \mathbf{b}_s) \quad (13)$$

$$\mathbf{p}_t = \text{softmax}(\mathbf{U}_p\mathbf{s}_t + \mathbf{b}_p) \quad (14)$$

Here s_t refers to the pre-softmax hidden representation with concatenated inputs $h_t, x_t, \phi_t(V)$ and p_t is the probability distribution over the vocabulary for sampling the current word.

4. Key Addressing

Soft attention mechanism have recently been successful in Image Captioning[45] and Video Description Generation [47] tasks because they help to focus on relevant parts of the features rather than mean pooling. In [41], the feature vectors were simply averaged, leading to loss of the temporal relationships between frames of the video. Soft attention mechanism weights each frame allowing to exploit the temporal structure in video. They were also used in the Key-Value Memory Networks[25] to focus on more relevant keys and reads the weighted sum of values.

We propose a recurrent key addressing mechanism which looks at the previous attention distribution over keys to obtain the new relevance scores. This helps to exploit multiple hierarchies in sequence-to-sequence architectures as the model is selecting relevant frames based on the key distribution seen so far and the previously generated words. There is a Key Addressing RNN (referred to as Memory LSTM in Fig. 4) to keep track of the previous attention distributions over keys and uses the hidden state to find the new attention weights α_i. The current hidden state of the Memory LSTM takes as input the previous relevance distribution over keys $\phi_{t-1}(K)$:

$$\mathbf{h}_t^k = \mathbf{f}^k(\phi_{t-1}(K), h_{t-1}^k) \quad (15)$$

where f^k is the recurrent unit.

The relevance distribution on the memory slots is evaluated using a query vector q, which summarizes the frames seen so far and the output generated. The query vector q is a weighted combination of the decoder and key-addressing hidden states.

$$q = W_k h_t^k + W_d h_{t-1} \quad (16)$$

For obtaining the attention weights, the relevance score e_i^t of i-th temporal feature is obtained using the decoder RNN hidden state h_{t-1}, key addressing RNN hidden state h_t^k and the i-th key vector k_i:

$$e_i^t = \mathbf{w}_t \tanh(\mathbf{q} + \mathbf{U}_a k_i) \quad (17)$$

where $\mathbf{w}_t$, $\mathbf{W}_d$, $\mathbf{W}_k$ and $\mathbf{U}_a$ are parameters of the model. This allows us take into consideration the previously generated words, the attention distribution on previous time steps and the individual key representations.

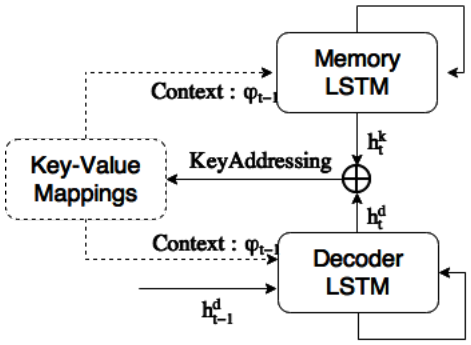

Figure 3. Key addressing LSTM is shown here. The memory LSTM updates its hidden state using the last attention distribution over keys. The new hidden state is used with decoder last hidden state to get new relevance scores which is combined with values and passed to the decoder to generate the next word.

These relevance scores are normalised using a softmax function which gives the attention distribution α_i^t using:

$$\alpha_i^t = \frac{exp\{e_i^t\}}{\sum_{j=1}^{N} exp\{e_j^t\}} \quad (18)$$

The proposed model is flexible, because it allows multiple hops for iterative inference just like in the original framework. The query vector q at initial time step is a mean-pooled average of all the keys k_i. The segregation of the vision and language components into key-value pairs provides a better context for the RNN decoder. Also, the explicit memory structure provides access to the image frames at all time steps allowing the model to assign weights to the key-frames without losing information.

5. Experimental Setup

5.1. Dataset

Youtube2Text The proposed approach is benchmarked on the Youtube2Text[7] dataset which consists of 1,970 Youtube videos with multiple descriptions annotated through Amazon Mechanical Turk. The videos are generally short (9 seconds on an average), and depict a single activity. Activities depicted are open domain ranging from everyday objects to animals, scenarios, actions, landscapes. etc. The dataset consists of 80,839 annotations and vocabulary of 16,000 words approximately, with an average of 41 annotations per clip and 8 words per sentence respectively. The training, validation and test sets have 1,200, 100 and 670 videos respectively which is exactly the same splits as in previous work on video captioning [47, 4, 27].

Key-Value Memories We select 28 equally spaced frames and pass them through a pretrained VGG-16[33] and GoogleNet[38] because of their state of the art performance in object detection on Imagenet[12] database.

Table 1. Experiment results on the Youtube2Text Dataset.

Model	BLEU@4	METEOR	CIDEr	Feat.	Fine
VGG-Encoder	0.404	0.295	0.515	No	No
GoogLeNet-Encoder	0.427	0.303	0.534	No	No
t-KeyAddressing (G)	0.436	0.308	0.545	No	No
m-KeyAddressing (Memory LSTM) (G)	**0.457**	**0.319**	**0.573**	No	No
Enc-Dec Basic(Yao et al. [47])	0.3869	0.2868	0.4478	No	No
GoogLeNet + HRNE(Pan et al. [27])	0.438	**0.321**	.	No	No
C3D(Yao et al. [47])	0.4192	0.2960.	0.5167	Yes	No
VGG + C3D + p-RNN(Yu et al.[49]	**0.499**	**0.326**	-	Yes	No
S2VT(Venugopalan et al. [40])	-	.298	-	Yes	No
GRU-RCN(Ballas et al. [4])	**0.490**	0.3075	**0.5937**	Yes	Yes

For an input image of size WXH, visual features with shape $(\lfloor \frac{W}{16} \rfloor, \lfloor \frac{W}{16} \rfloor, C)$ with C as 512 are extracted from the *conv5_3* layer of VGG-16. We simply average over a feature map which results in a feature of length C and are then used as keys. The visual features extracted from the *pool5/7x7_s1* layer of GoogLeNet is a 1024 dimensional vector and are also used as keys.

The values are generated from a pre-trained image captioning module following [45, 16], which identifies salient regions in an image, and generates a caption for each of these regions. Densecap jointly models the task of object localization as well as description using a Fully Convolutional Localization Network (FCLN). After passing the image through CNN and Localization layers, Recognition network is used which is essentially a fully connected layer. We extract output from this layer which is encoded as region codes of size BxD, where B is the number of salient regions or boxes, and D is the representation with dimension 4096. Along with the features, a score S is assigned to each of the regions which denotes its confidence. A weighted sum of features of top 5 scores is calculated to get values with dimension D.

Preprocessing: The video descriptions are tokenized using the wordpunct_tokenizer from the NLTK toolbox[24]. The number of unique words were 15,903 in the Youtube2Text dataset.

5.2. Model Specifications

We vary the underlying CNN for feature encoding, and conduct ablation studies with key addressing with and without LSTM which results in four model variations. The results from different variations are presented in Table 1. *VGG-Encoder* uses features encoded from the last convolution layer in VGG-16 CNN [33] along with a LSTM decoder. Here, we define $q = h_{t-1}$ for the key addressing following [25]. *GoogLeNet-Encoder* is exactly the same as *VGG Encoder* with the only difference being GoogLeNet used for feature encoding. **t-KeyAddressing** adds to *GoogLeNet-Encoder* by addressing keys through both the last hidden state h_{t-1} and the previous time-step attention distribution on the keys. Finally, **m-KeyAddressing** summarizes all the previous attention distributions through a LSTM, after which its hidden state is used for key addressing along with last hidden state of the Decoder. These variations help us identify architecture changes that lead to large improvements on the evaluation metric. Detailed comparison of the results is provided in Section 5.6.

5.3. Model Comparisons

We compare the model performance with previous state of the art approaches and some strong baselines. Pan et al. [28] explicitly learn a visual-semantic joint embedding model for exploiting the relationship between visual features and generated language, which is then used in a similar encoder-decoder framework. Yao et al.[47] utilizes a temporal attention mechanism for global attention apart from local attention using 3-D Convolution Networks. Ballas et al.[4] proposed an encoder to learn spatial-temporal features across frames, introducing a variant GRU with convolution operations (GRU-RCN). In the current state-of-art Yu et al. [49] model the decoder as a paragraph generator, describing the videos over multiple sentences using stacked LSTMs.

5.4. Evaluation Metrics

We evaluate our approach using standard evaluation metrics for video captioning tasks, namely BLEU [29], METEOR [22] and CIDEr[39] to compare the generated sequences with the human annotations. Treating this as a Sequence-to-Sequence model, and comparing against given ground truth, higher scores on the metrics indicate better performance. The metrics score are calculated based on alignment and similarity between the generated and candidate reference descriptions. We use the code accompanying the Microsoft COCO Evaluation script [8] to obtain the results reported in the paper.

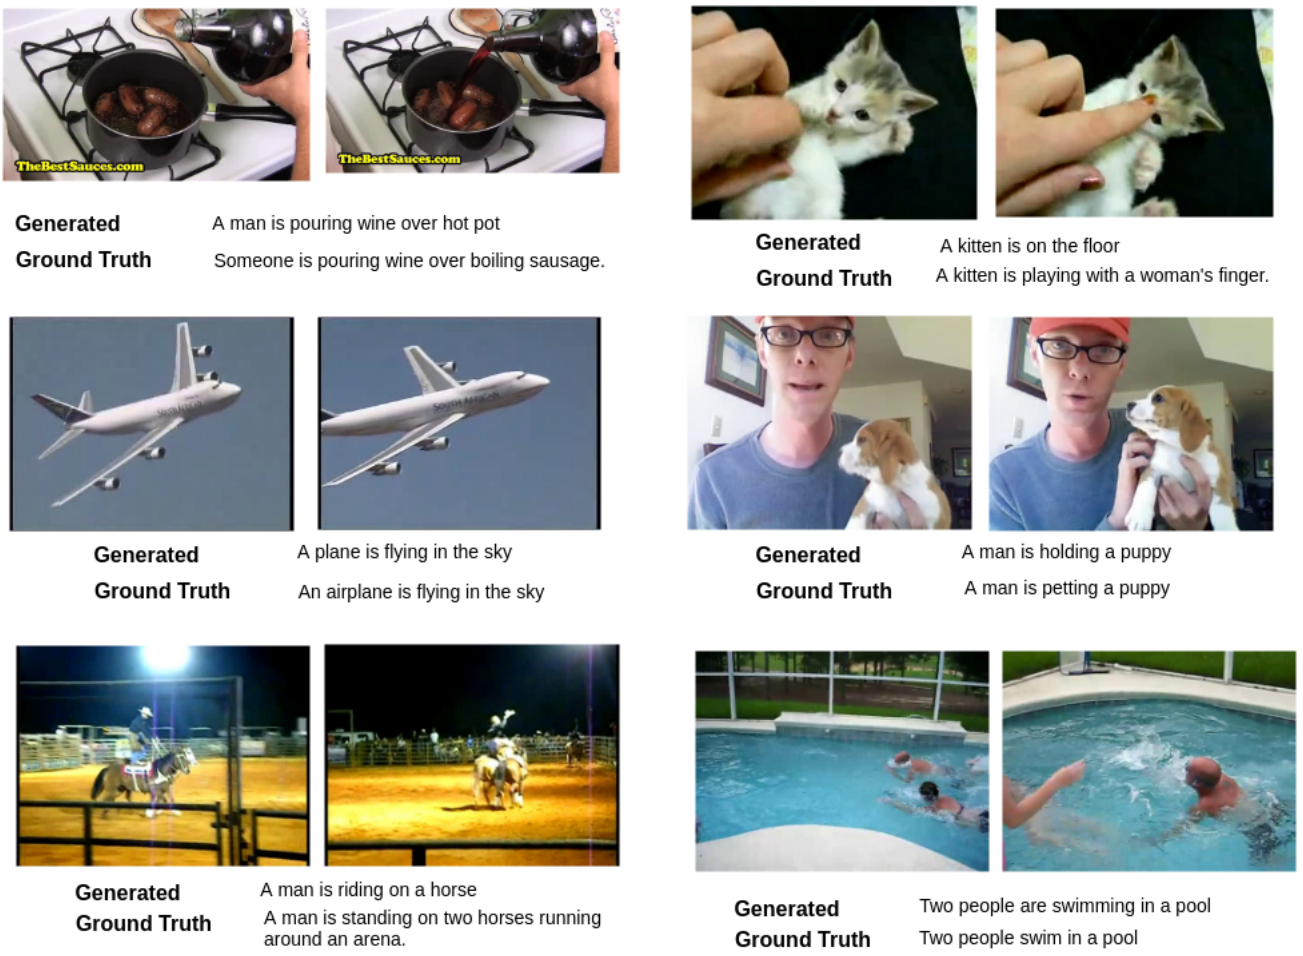
Generated
A man is pouring wine over hot pot
Ground Truth
Someone is pouring wine over boiling sausage.
Generated
A kitten is on the floor
Ground Truth
A kitten is playing with a woman's finger.
Generated
A plane is flying in the sky
Ground Truth
An airplane is flying in the sky
Generated
A man is holding a puppy
Ground Truth
A man is petting a puppy
Generated
A man is riding on a horse
Ground Truth
A man is standing on two horses running around an arena.
Generated
Two people are swimming in a pool
Ground Truth
Two people swim in a pool

Figure 4. Sample results generated by our model on the Youtube2Text dataset.

5.5. Training Details

The model predicts the next output word given the previous words and the input video. Thus, the goal is to maximize the log likelihood of the loss function:

$$L = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{|y^i|}\log p(y_j^i|y_{<j}^i, \mathbf{x}^n, \theta) \qquad (19)$$

where N is the total number of video-description pairs and length of each description y_i is $|y_i|$. Here x^n refers to the input video provided as context to the decoder. We train our network parameters θ through first order stochastic gradient-based optimization with an adaptive learning rate using the Adadelta [51] optimizer. We set the batch size to be 64 in our implementation. We optimize hyperparameters, namely number of hidden units in Decoder LSTM, key addressing LSTM, learning rate, word embedding dimension for the log loss using random search [6].

5.6. Results

Table 5.1 summarizes the performance of our approach, and compares it with various state of the art techniques. We test our proposed model in the basic setting with visual features extracted from a pretrained VGG-16 [33] and GoogLeNet[38] models, as depicted by the first two lines of Table 1. The key addressing in these models, follows the basic approach as in [25], i.e without tracking the last time steps. These models are able to outperform the baselines of S2VT [40], and the Basic Encoder-Decoder model introduced in [47] on all three metrics. Using the recurrent key-addressing, we see further boost in performance. We observe that using features from pretrained GoogleNet rather than VGG-16 improves on the results, following [42].

t-KeyAddressing refers to using only the last time attention distribution over keys in the query vector. Using GoogleNet to extract features, we achieve significantly better results compared to the temporal attention introduced by [47], where they do not utilize attention distribution from previous time steps further in the model. Additionally, [47]

use 3D Convolution Networks to extract local temporal features, while we work directly on the individual frame features.

Using a recurrent addressing scheme; an LSTM for storing previous attention distributions, leads to huge improvement on all the metrics and demonstrates the effectiveness of the proposed framework. This model, referred to as **m-KeyAddressing**, outperforms strong baselines from [27] by a significant margin on BLEU@4. While the improvements on METEOR are significant compared to **t-KeyAddressing**, [27] learns a hierarchical representation of the video, thereby performing slightly better. In the current setting, our model is unable to outperform [49] which uses a stacked LSTM at the decoder and generates more granular descriptions. A possible approach for improvement in this direction is using more sophisticated regularizers for training the decoder, as proposed in [2, 21].

In Table 5.1 we provide comparison on whether the models use finetuning on the CNN encoder (represented by *Fine*) or if they use external features, like on action recognition, optical flow (represented by *Feat*). It is to be noted, that we do not finetune the encoder, or the image captioning module as compared to [4] which finetunes the encoder CNN on UCF101 action recognition set[34]. Also, no additional features are extracted for gaining more information about motion, actions etc. as in [49], [4], [40].

In Fig 4, we show examples of some of the input frames and generated outputs, with ground truths. Some of the examples demonstrate the fact that the model is able to infer the activities from the video frames, like "swimming", "riding" and "flying" which is often distributed across multiple frames.

6. Conclusion

We introduce a recurrent memory addressing model for video captioning. By decomposing the visual and language components, we explicitly exploit temporal structure a multiple hierarchies. Extensive experiments on the proposed model outperform strong baselines across several metrics, and achieve competitive scores against the state-of-the-art benchmarks. To the best of our knowledge this is the first proposed work for video-captioning in a Memory Networks setting, and does not rely on heavily annotated videos to generate intermediate semantic-embedding for supporting the decoder. Further work would be exploring the effectiveness of the model on longer videos, and generating fine-grained descriptions with more sophisticated decoders.

References

[1] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. Lawrence Zitnick, and D. Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433, 2015. 1

[2] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 8

[3] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014. 1, 2, 3

[4] N. Ballas, L. Yao, C. Pal, and A. Courville. Delving deeper into convolutional networks for learning video representations. *arXiv preprint arXiv:1511.06432*, 2015. 2, 4, 5, 6, 8

[5] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994. 4

[6] J. Bergstra and Y. Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb):281–305, 2012. 7

[7] D. L. Chen and W. B. Dolan. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 190–200. Association for Computational Linguistics, 2011. 2, 5

[8] X. Chen, H. Fang, T.-Y. Lin, R. Vedantam, S. Gupta, P. Dollár, and C. L. Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 6

[9] K. Cho, B. Van Merriënboer, Ç. Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar, Oct. 2014. Association for Computational Linguistics. 2, 3

[10] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014. 1

[11] P. Das, C. Xu, R. F. Doell, and J. J. Corso. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2634–2641, 2013. 2

[12] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009. 5

[13] H. Fang, S. Gupta, F. Iandola, R. K. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. C. Platt, et al. From captions to visual concepts and back. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1473–1482, 2015. 2

[14] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov, et al. Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, pages 2121–2129, 2013. 2

[15] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 1, 3, 4

[16] J. Johnson, A. Karpathy, and L. Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. *arXiv preprint arXiv:1511.07571*, 2015. 1, 2, 4, 6

[17] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3128–3137, 2015. 3

[18] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014. 2

[19] G. Kulkarni, V. Premraj, V. Ordonez, S. Dhar, S. Li, Y. Choi, A. C. Berg, and T. L. Berg. Babytalk: Understanding and generating simple image descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12):2891–2903, 2013. 2

[20] A. Kumar, O. Irsoy, J. Su, J. Bradbury, R. English, B. Pierce, P. Ondruska, I. Gulrajani, and R. Socher. Ask me anything: Dynamic memory networks for natural language processing. *arXiv preprint arXiv:1506.07285*, 2015. 3

[21] A. Lamb, A. Goyal, Y. Zhang, S. Zhang, A. Courville, and Y. Bengio. Professor Forcing: A New Algorithm for Training Recurrent Networks. *ArXiv e-prints*, Oct. 2016. 8

[22] M. D. A. Lavie. Meteor universal: language specific translation evaluation for any target language. *ACL 2014*, page 376, 2014. 6

[23] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 1

[24] E. Loper and S. Bird. Nltk: The natural language toolkit. In *Proceedings of the ACL-02 Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics - Volume 1*, ETMTNLP '02, pages 63–70, Stroudsburg, PA, USA, 2002. Association for Computational Linguistics. 6

[25] A. Miller, A. Fisch, J. Dodge, A.-H. Karimi, A. Bordes, and J. Weston. Key-value memory networks for directly reading documents. *arXiv preprint arXiv:1606.03126*, 2016. 2, 3, 4, 5, 6, 7

[26] V. Mnih, N. Heess, A. Graves, et al. Recurrent models of visual attention. In *Advances in Neural Information Processing Systems*, pages 2204–2212, 2014. 2

[27] P. Pan, Z. Xu, Y. Yang, F. Wu, and Y. Zhuang. Hierarchical recurrent neural encoder for video representation with application to captioning. *arXiv preprint arXiv:1511.03476*, 2015. 2, 5, 6, 8

[28] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui. Jointly modeling embedding and translation to bridge video and language. *arXiv preprint arXiv:1505.01861*, 2015. 2, 6

[29] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics, 2002. 6

[30] T. Rocktäschel, E. Grefenstette, K. M. Hermann, T. Kočiskỳ, and P. Blunsom. Reasoning about entailment with neural attention. *arXiv preprint arXiv:1509.06664*, 2015. 3

[31] M. Rohrbach, W. Qiu, I. Titov, S. Thater, M. Pinkal, and B. Schiele. Translating video content to natural language descriptions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 433–440, 2013. 2

[32] I. V. Serban, A. Sordoni, Y. Bengio, A. Courville, and J. Pineau. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16)*, 2016. 1

[33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2, 3, 5, 6, 7

[34] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 8

[35] M. F. Stollenga, J. Masci, F. Gomez, and J. Schmidhuber. Deep networks with internal selective attention through feedback connections. In *Advances in Neural Information Processing Systems*, pages 3545–3553, 2014. 1

[36] S. Sukhbaatar, J. Weston, R. Fergus, et al. End-to-end memory networks. In *Advances in neural information processing systems*, pages 2440–2448, 2015. 3

[37] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014. 2

[38] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015. 2, 5, 7

[39] R. Vedantam, C. Lawrence Zitnick, and D. Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4566–4575, 2015. 6

[40] S. Venugopalan, M. Rohrbach, J. Donahue, R. Mooney, T. Darrell, and K. Saenko. Sequence to sequence – video to text. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 2, 6, 7, 8

[41] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko. Translating videos to natural language using deep recurrent neural networks. In *NAACL HLT*, 2015. 2, 5

[42] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3156–3164, 2015. 2, 3, 4, 7

[43] P. J. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990. 2

[44] J. Weston, S. Chopra, and A. Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014. 3

[45] K. Xu, J. Ba, R. Kiros, K. Cho, A. C. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. *CoRR*, abs/1502.03044, 2015. 1, 2, 3, 4, 5, 6

[46] Z. Yang, Y. Yuan, Y. Wu, R. Salakhutdinov, and W. W. Cohen. Encode, review, and decode: Reviewer module for caption generation. *arXiv preprint arXiv:1605.07912*, 2016. 2, 3

[47] L. Yao, A. Torabi, K. Cho, N. Ballas, C. Pal, H. Larochelle, and A. Courville. Describing videos by exploiting temporal

structure. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4507–4515, 2015. 1, 2, 5, 6, 7

[48] Q. You, H. Jin, Z. Wang, C. Fang, and J. Luo. Image captioning with semantic attention. *arXiv preprint arXiv:1603.03925*, 2016. 3, 4

[49] H. Yu, J. Wang, Z. Huang, Y. Yang, and W. Xu. Video paragraph captioning using hierarchical recurrent neural networks. *arXiv preprint arXiv:1510.07712*, 2015. 2, 6, 8

[50] W. Zaremba, I. Sutskever, and O. Vinyals. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*, 2014. 3

[51] M. D. Zeiler. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012. 7